

Computer Networks

ICS 651

- packet networking
- Internet Protocol
- routing and project1
- exam

Packet Networking

- sharing an expensive medium
- assumes nobody needs to send all the time (is it correct?)
- shorter packets mean shorter delay for small amounts of information
- bandwidth sharing mostly implicit, though many attempts to make it more explicit, manageable, controllable

Packet Networking, continued

- a very simple system can be built on top of point-to-point connections which can transfer one byte at a time
- alternately, shared media have conventions (Medium Access Control, MAC) for transferring entire packets and controlling access to a shared medium
- an incoming packet must get from the hardware to the application: downcall or upcall?
- sending a packet only requires queueing if the medium is busy

DNS

- structured, hierarchical name space
- arbitrary mapping of domain names to IP addresses
- mapping supported by distributed database
- different "types" of mapping, e.g. A for address, MX for mail server, NS for Name Server
- recursive or iterative queries
- queries elicit answers, which may include authority and additional information
- idempotent protocol means queries can be repeated for reliability

Internet Protocol

- end-to-end addressing, protocol, length, fragmentation
- hop-by-hop TTL/HL, checksum for IPv4
- best-effort: packets can be discarded, reordered, misdelivered, delayed, or have bit errors
- once the routing tables are in place, forwarding is straightforward:
 - verify header checksum, discard if incorrect
 - deliver to socket if destination is local
 - decrement TTL/HL, discard if zero
 - look up destination in routing table, select (a) longest match, (b) if more than one longest match, longest match with best metric, or (c) if more than one of b., first matching entry
 - recompute checksum
 - queue on output interface

Internet Protocol, continued

- network masks allow the division of networks into subnetworks
- all-1s host part is broadcast address, all-0s host part is network number
 - e.g. for address 1.2.3.4, netmask 255.255.0.0, network number is 1.2.0.0, broadcast address is 1.2.255.255
- 255.255.192.0 is a netmask with how many bits?
- ICMP used to report errors (e.g. when TTL/HL=0 during traceroute) or probe network (ping)
- IPv6: much bigger addresses, fewer processing steps at router
 - especially, no recomputation of checksum
- path MTU discovery
- fragmentation: why needed, packet identifiers, fragment offsets

Routing and Distance Vector

- Routing: exchanging information to build routing tables
- Distance-Vector: acquire routing tables from neighbors, use in computing own routes
- RIP, RIPv2
- BGP uses a variant, path-vector, that allows more administrative control
- information received from neighbors may reflect old state that is no longer correct
- split-horizon, with or without poisoned reverse, helps prevent some of the routing loops
- other link failures lead to counting-to-infinity

Routing, continued

- Link-State: the networks that each router is connected to, are broadcast to all other routers
- routers use Dijkstra's shortest path algorithm to compute the best route to each destination
- OSPF, IS-IS use LS
- OSPF divides the AS into areas
- routers in each area receive complete information about the other routers in the same area, and summary information about routers in other areas
- OSPF, IS-IS, and RIP are Interior Gateway Protocols, IGPs, and only work within an Autonomous System
- BGP is an EGP, Exterior Gateway Protocol

Project 1

- simulated networks
- unreliable transmission
- upcalls and threads
- basic IP forwarding
- distance-vector routing
- homework 1: sockets API
- homework 2: UDP and simulated network

Exam

- open book, open notes, open Internet
 - closed friends!!!
- answer the question that is asked
- answer each question completely if possible
 - incomplete answers get partial credit, so do what you can
- essay questions
 - maybe simple programming questions
- can be done during class time only:
 - mailed at the beginning of class,
 - due at the end of class
 - best if mailed 1-5 minutes early!